arXiv:1703.09975v1 [stat.ML] 29 Mar 2017

Improving Spectral Clustering using the Asymptotic Value of the Normalised Cut

David P. Hofmeyr*
Department of Statistics and Actuarial Science, Stellenbosch University

March 30, 2017

Abstract

Spectral clustering is a popular and versatile clustering method based on a relaxation of the normalised graph cut objective. Despite its popularity, however, there is no single agreed upon method for tuning the important scaling parameter, nor for determining automatically the number of clusters to extract. Popular heuristics exist, but corresponding theoretical results are scarce. In this paper we investigate the asymptotic value of the normalised cut for an increasing sample assumed to arise from an underlying probability distribution, and based on this result provide recommendations for improving spectral clustering methodology. A corresponding algorithm is proposed with strong empirical performance.

Keywords: Automatic clustering, Low density separation, Self-tuning clustering

1 Introduction

Clustering is the problem of identifying relatively homogeneous groups (clusters) within a given dataset. Without a universal definition of what constitutes a cluster, multiple objectives have been proposed and a multitude of algorithms developed for each objective. A common and intuitive theme underlying most approaches is that data within the same cluster should be more similar than data in different clusters. A principled approach to addressing this relies on the construction of a similarity graph for the data, in which vertices correspond to data and edges are assigned a weight equal the similarity between their vertex endpoints. A graph cut refers to the removal of a subset of edges from the graph which renders it disconnected. A cut of the similarity graph for which the edges removed have low total weight therefore induces a partition (a clustering) of the data for which the total similarity between data in different clusters is low. Such a cut, however, does not ensure that the similarity of data in the same cluster is comparatively high. Moreover, in practice the minimum weight cut tends to separate only small collections of points from the remainder, which may not constitute complete clusters (von Luxburg, 2007). To account for these facts, normalisations of the graph cut objective have been proposed (Shi and Malik, 2000; Hagen and Kahng, 1992). Normalisation, however, renders the associated optimisation problem NP-hard (Wagner and Wagner, 1993). It has been shown that a continuous relaxation of the normalised graph cut problem can be solved using the eigenvectors of the so-called graph Laplacian matrix. It is this relaxation of the normalised graph cut problem which is given the name spectral clustering.

Spectral clustering has become extremely popular in recent years, for its versatility and strong performance in numerous application areas (Ng et al., 2001; von Luxburg, 2007). One of the appealing properties of spectral clustering is its ability to identify highly non-convex clusters. Of crucial importance to the success of spectral clustering, however, is an appropriate measure of similarity to be used in the construction of the similarity graph. In addition, as with all clustering methods, determining automatically the number of clusters to extract from the data is a non-trivial task. In general the similarity between two points is determined by a decreasing function of the distance between them, so that pairs of data which are nearer in space tend to be assigned higher similarity than those which are more distant. For data belonging to $\mathbb{R}^d$, it is most common to define the similarity between $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ as $k(\|\mathbf{x}-\mathbf{y}\|/\sigma)$, where $k : \mathbb{R}^+ \to \mathbb{R}^+$ is a kernel function,

*The author would like to thank Dr. Nicos Pavlidis and Ms. Katie Yates for their assistance with experiments, and Dr. Nicos Pavlidis for his valuable comments on this work.

$\|\cdot\|$ is the Euclidean norm, and the parameter $\sigma > 0$ captures the scale of the data. The crucial factor then becomes the setting of the scaling parameter, σ. There is no single agreed upon method for setting this parameter. Some popular heuristics have been developed (Ng et al., 2001; Zelnik-Manor and Perona, 2004), but few of these are supported by theory. If σ is set too large, then the ability of spectral clustering to separate highly non-convex clusters is severely diminished. On the other hand, if σ is set too small, then spectral clustering becomes highly susceptible to separating only outliers rather than complete clusters of data. Existing approaches for selecting the number of clusters rely either on the eigen-gaps of the graph Laplacian, or on the linear structure of the associated eigenvectors (Zelnik-Manor and Perona, 2004; von Luxburg, 2007).

In this paper we will use the asymptotic properties of the normalised cut to provide recommendations for improving spectral clustering methodology. Specifically, it will be shown that if the data are assumed to represent a sample of realisations of a continuous random variable with Lipschitz continuous density, then the (scaled) normalised cut measured across a smooth surface converges almost surely to the normalised integrated squared density over that surface. This is a specific case of the result of Trillos et al. (2015), for which we provide an alternative and what we believe to be more accessible proof. A similar investigation was conducted by Narayanan et al. (2006), however our conclusions are substantially different. Relating normalised cuts to the probability distribution underlying the data helps to better understand the normalised cut objective, and therefore spectral clustering. This allows us to provide sound and justified recommendations for the practical implementation of spectral clustering methods. Importantly, it will be discussed how the minimum normalised cut solution tends to separate clusters by regions of low density, while at the same time separating high density regions in the underlying density function. These are key features in the non-parametric statistical approach to clustering, known as *density clustering*, in which clusters are associated with regions of high probability density surrounding the modes of the density function. This result is especially appealing as the density clustering objective provides a principled means for automatically determining the number of clusters; as the number of modes of the density function. Furthermore, the conditions on the scaling parameter, σ, required for convergence provide guidelines for establishing an order of magnitude for this parameter, and also dictate how it should decrease as sample sizes increase. Finally, this result provides some insight into the susceptibility of spectral clustering to outliers.

Based on these observations we propose a spectral clustering algorithm which we have found to exhibit strong performance in a range of applications. We are not aware of any existing algorithms which make use of the asymptotic properties of the normalised cut objective to improve clustering methodology. This is also the only spectral clustering algorithm of which we are aware which uses the structure of the clusters in the original data space, and not the eigenvector representation, to select the number of clusters.

The remainder of this paper is organised as follows. In Section 2 a background on normalised cuts and spectral clustering will be provided. Following that in Section 3 we will investigate the asymptotic value of the normalised cut, and discuss in greater detail the practical implications of this result. We also propose an algorithm which makes use of these practical implications to automatically set the scaling parameter and select the number of clusters to extract. The results of experiments on publicly available benchmark datasets are then presented in Section 4. A final discussion is provided in Section 5.

2 Normalised Cuts and Spectral Clustering

In this section we provide a brief introduction to normalised cuts and the relaxed solution through spectral clustering. Note that the spectral clustering methodology is applicable to any data for which pair-wise similarities can be established, however we will focus only on Euclidean embedded data for this work. Let $\mathcal{X} = \{\mathbf{x}_1, ..., \mathbf{x}_n\}$ be a collection of data in $\mathbb{R}^d$. For each pair $\mathbf{x}_i, \mathbf{x}_j$ define their similarity as $k(\|\mathbf{x}_i - \mathbf{x}_j\|/\sigma)$, where $k : \mathbb{R}^+ \to \mathbb{R}^+$ is a kernel, and $\sigma > 0$ is a scaling parameter. A popular choice for k is the Gaussian kernel, $k(x) = \exp(-x^2/2)$. Now, define the *affinity matrix* $A \in \mathbb{R}^{n\times n}$ s.t.

$$A_{i,j} = \begin{cases} k(\|\mathbf{x}_i - \mathbf{x}_j\|/\sigma), & i \neq j \\ 0, & i = j, \end{cases} \tag{1}$$

and the corresponding *degree matrix* $D = \text{diag}\left(\sum_{j=1}^n A_{1,j}, ..., \sum_{j=1}^n A_{n,j}\right)$. The matrix A uniquely defines a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the similarity graph of the data, which has n vertices (one per datum) and whose edges have weights given by the elements of A, i.e. $\mathcal{E}_{i,j} = A_{i,j}$.

Remark 1. *It is important to note that some authors prefer to set the diagonal elements of A to $k(0) = 1$. This is a matter of preference, and we have found that setting the diagonal of A to zero improves performance. The asymptotic result we will show in the next section applies for both definitions of A.*

Now, for a partition (clustering) of $\mathcal{X}$ into $\{\mathcal{C}_1, ..., \mathcal{C}_c\}$, the corresponding graph cut is given by,

$$\text{Cut}(\mathcal{C}_1, ..., \mathcal{C}_c) = \frac{1}{2} \sum_{i=1}^{c} \sum_{\substack{j,l: \mathbf{x}_j \in \mathcal{C}_i, \\ \mathbf{x}_l \notin \mathcal{C}_i}} A_{j,l}. \tag{2}$$

The factor $\frac{1}{2}$ is in place since each edge removed to create the cut is counted twice in the above summation. Notice that the total cut for such a partition can be re-expressed as $\sum_{i=1}^{c} \text{Cut}(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i)$. A low value cut ensures that the total similarity between data in different clusters is low, however it does not necessarily result in clusters with high internal similarity. A preferable objective for clustering is the *normalised cut* (Shi and Malik, 2000), defined as

$$\text{NCut}(\mathcal{C}_1, ..., \mathcal{C}_c\} = \frac{1}{2} \sum_{i=1}^{c} \frac{1}{\sum_{\mathbf{x}_j \in \mathcal{C}_i} D_{j,j}} \sum_{\substack{j,l: \mathbf{x}_j \in \mathcal{C}_i \\ \mathbf{x}_l \notin \mathcal{C}_i}} A_{j,l} = \sum_{i=1}^{c} \frac{\text{Cut}(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i)}{\sum_{\mathbf{x}_j \in \mathcal{C}_i} D_{j,j}}. \tag{3}$$

The quantity $\sum_{\mathbf{x}_j \in \mathcal{C}} D_{j,j}$, where $\mathcal{C} \subset \mathcal{X}$, is referred to as the *volume* of $\mathcal{C}$, denoted $\text{vol}(\mathcal{C})$ (von Luxburg, 2007). Notice that a low value normalised cut will, in general, have a low value of $\text{Cut}(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i)$ and a large value of $\text{vol}(\mathcal{C}_i)$, for each $i = 1, \ldots, c$. Now, recall that $D_{j,j} = \sum_{l=1}^{n} A_{j,l}$, and therefore,

$$\begin{aligned}
\text{vol}(\mathcal{C}_i) &= \sum_{\substack{j,l: \mathbf{x}_j \in \mathcal{C}_i \\ \mathbf{x}_l \in \mathcal{X}}} A_{j,l} \\
&= \sum_{\substack{j,l: \mathbf{x}_j \in \mathcal{C}_i \\ \mathbf{x}_l \notin \mathcal{C}_i}} A_{j,l} + \sum_{j,l: \mathbf{x}_j, \mathbf{x}_l \in \mathcal{C}_i} A_{j,l} \\
&= \text{Cut}(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i) + \sum_{j,l: \mathbf{x}_j, \mathbf{x}_l \in \mathcal{C}_i} A_{j,l}.
\end{aligned}$$

For $\text{Cut}(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i)$ to be low and $\text{vol}(\mathcal{C}_i)$ to be high, we must therefore have $\sum_{j,l: \mathbf{x}_j, \mathbf{x}_l \in \mathcal{C}_i} A_{j,l}$ being relatively high. This quantity represents the internal cluster similarity for cluster $\mathcal{C}_i$. The minimum normalised cut therefore tends to produce clustering solutions which have low between cluster similarity, and high within cluster similarity.

It has been shown that determining the optimal partition of $\mathcal{X}$ based on the normalised cut objective is NP-hard (Wagner and Wagner, 1993). However, a reformulation of the normalised cut objective, in terms of the so-called *graph Laplacian matrix*, leads to an intuitive relaxation of the problem (Shi and Malik, 2000; von Luxburg, 2007). The graph Laplacian may be defined as,

$$L = D^{-1/2}(D - A)D^{-1/2} = I_n - D^{-1/2}AD^{-1/2}. \tag{4}$$

Now, for a partition of $\mathcal{X}$, as before, into $\{\mathcal{C}_1, \ldots, \mathcal{C}_c\}$, define the matrix

$$V \in \mathbb{R}^{n \times c} : V_{i,j} = \begin{cases} \sqrt{1/\text{vol}(\mathcal{C}_j)}, & \mathbf{x}_i \in \mathcal{C}_j \\ 0, & \mathbf{x}_i \notin \mathcal{C}_j. \end{cases} \tag{5}$$

Then the minimum normalised cut problem for a fixed number of clusters, c, can be restated as (von Luxburg, 2007),

$$\begin{aligned}
&\min_{\mathcal{C}_1, ..., \mathcal{C}_c} \text{trace}(V^\top (D - A)V) \\
&\quad \text{s.t. } V \text{ defined as in Eq. (5)} \\
&\quad V^\top DV = I_c.
\end{aligned} \tag{6}$$

Relaxing the discreteness condition on the elements of V in Eq. (5) and setting $U = D^{1/2}V$ we arrive at

$$\min_{U \in \mathbb{R}^{n \times c}} \text{trace}(U^\top LU), \text{ s.t. } U^\top U = I_c, \tag{7}$$

which is the canonical formulation of the problem of finding the eigenvectors of L associated with its smallest c eigenvalues. The approximate solution of the minimum normalised cut problem is therefore given by $D^{-1/2}U$, where $U \in \mathbb{R}^{n \times c}$ has as columns the c eigenvectors of L associated with its c smallest eigenvalues.

Notice that the optimal V in problem (6) has as columns scaled cluster indicator vectors for the clusters $\mathcal{C}_1, \ldots, \mathcal{C}_c$, and hence the cluster assignment based on V is immediate. The relaxed solution $D^{-1/2}U$ arising from problem (7) does not necessarily take only discrete values. To obtain a final clustering solution, a simple clustering algorithm is applied to the rows of $D^{-1/2}U$. A popular choice for this final step is the k-means algorithm.

3 Improving Spectral Clustering using the Asymptotic Value of the Normalised Cut

In Section 2 the values of σ, the scale parameter, and c, the number of clusters, were assumed known. In practice determining these parameters is extremely challenging. In this section we will investigate the asymptotic value of the normalised cut for an increasing sample assumed to arise from an underlying probability distribution. Specifically, we show that the asymptotic value of the normalised cut measured across a smooth surface converges almost surely to the normalised integrated squared density measured along that surface. This allows us to provide recommendations for automatically setting values of σ and c. Based on these recommendations we propose an algorithm which we have found to exhibit strong empirical performance in a range of applications.

3.1 The asymptotic value of the normalised cut

The result presented here is a specific case of the result of Trillos et al. (2015); where we investigate the particular practical example of the normalised cut objective, using the Gaussian similarity kernel. By investigating this particular case we are able to provide a more accessible proof. A similar investigation into the asymptotic value of the normalised cut is given in Narayanan et al. (2006), however the result we obtain herein is both more intuitive in relation to normalised cuts and is a preferable objective for clustering.

Theorem 2. *Let $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ be an i.i.d. sample of realisations of a continuous random variable, $\mathbf{X}$, over $\mathbb{R}^d$, with bounded and Lipschitz continuous probability density $p : \mathbb{R}^d \to \mathbb{R}^+$. Let S be a smooth surface which partitions $\mathbb{R}^d$ into S_1 and S_2 such that $0 < \mathbb{P}(\mathbf{X} \in S_1) < 1$. Let $\{\sigma_n\}_{n=1}^{\infty}$ be a null sequence satisfying $n\sigma_n^{2d+2+\epsilon} \to \infty$ for some $\epsilon > 0$. Then,*

$$\frac{\sqrt{2\pi}}{\sigma_n}\mathrm{NCut}(\mathcal{X} \cap S_1, \mathcal{X} \cap S_2) \xrightarrow{a.s.} \oint_S p(\mathbf{y})^2 d\mathbf{y} \left(\frac{1}{\int_{S_1} p(\mathbf{y})^2 d\mathbf{y}} + \frac{1}{\int_{S_2} p(\mathbf{y})^2 d\mathbf{y}} \right),$$

as $n \to \infty$.

We use $\oint$ to distinguish integrals over $d-1$ dimensional surfaces lying in $\mathbb{R}^d$ from regular integrals over subsets of $\mathbb{R}^d$ or $\mathbb{R}$, for which we use $\int$.

Remark 3. *The result in Theorem 2 can be shown to hold if the density, p, is only locally Lipschitz on a neigbourhood of S, and need not be globally Lipschitz. For ease of exposition, we prefer the above presentation.*

Notice that if we were able to maximise the minimum of the normalisation terms, $\int_{S_1} p(\mathbf{y})^2 d\mathbf{y}$ and $\int_{S_2} p(\mathbf{y})^2 d\mathbf{y}$, we would favour solutions in which there are regions within both S_1 and S_2 wherein the density takes values with especially large magnitude. This is because in maximising an integrated *squared* function, a solution in which the function takes large values, even over relatively small regions, would be preferred over a solution in which the function is of moderate magnitude over a relatively larger region. We expect, therefore, that the surface which minimises the normalised cut will (i) have low density along it, captured by the surface integral $\oint_S p(\mathbf{y})^2 d\mathbf{y}$, and (ii) separate regions each containing points at which p is large, as a result of normalising by the integrated squared density over S_1 and S_2. Specifically, if p is multimodal then we would expect that the surface which minimises the normalised cut will tend to both pass through low density regions, and also separate the modes of p. Both of these are key features in density clustering, in which clusters are associated with connected regions of high density surrounding the modes of the density, and are thus separated by regions of low density.

The proof of Theorem 2 is formed of the following two lemmas. Lemma 4 shows that the volume of the points arising within a measurable set, appropriately scaled, converges to the integrated squared density over that set. The intuition underlying this result is that for each $\mathbf{x}_i$, $\sum_{j=1}^{n} k(\|\mathbf{x}_i - \mathbf{x}_j\|/\sigma)$ is proportional to the kernel estimate of the density at $\mathbf{x}_i$, which we henceforth denote $\hat{p}(\mathbf{x}_i)$. The volume of the points arising in a measurable set, M, given by $\sum_{i:\mathbf{x}_i \in M} \sum_{j \neq i} k(\|\mathbf{x}_i - \mathbf{x}_j\|/\sigma)$, is therefore approximately proportional to a Monte Carlo estimate of $\int_M \hat{p}(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \approx \int_M p(\mathbf{x})^2 d\mathbf{x}$. Lemma 5 then shows that the (scaled) cut across a smooth surface converges to the integrated squared density measured across it. The proof is more technical, and involves a series of approximations. Some useful ideas are borrowed from Narayanan et al. (2006). The intuition in this case is that as the number of data increases, and so σ_n approaches zero, the only kernel weights with a non-negligible contribution to the cut are those lying close to the surface, S. Each point on S is therefore weighted proportionally to the number of pairs of data $\mathbf{x} \in S_1, \mathbf{y} \in S_2$ which have non-negligible similarity asymptotically. This weight is therefore proportional to squared density at the corresponding point on S. The complete proofs of Lemmas 4 and 5 are given in the appendix.

Lemma 4. *Let $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ be an i.i.d. sample of realisations of a continuous random variable, $\mathbf{X}$, over $\mathbb{R}^d$, with Lipschitz continuous probability density $p : \mathbb{R}^d \to \mathbb{R}^+$. Let $\{\sigma_i\}_{i=1}^{\infty}$ be a null sequence satisfying, $n\sigma_n^{2d+2+\epsilon} \to \infty$ for some $\epsilon > 0$. Let $M \subset \mathbb{R}^d$ be measurable. Then,*

$$\frac{c_d}{n^2 \sigma_n^d} \sum_{i:\mathbf{x}_i \in M} \sum_{j \neq i}^{n} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) \xrightarrow{a.s.} \int_M p(\mathbf{x})^2 d\mathbf{x}, \text{ as } n \to \infty,$$

where $c_d = (\int_{\mathbb{R}^d} k(\|\mathbf{x}\|) d\mathbf{x})^{-1}$.

Note that the rate of convergence of σ_n to zero required for the above result is actually less strict than stated. In addition the choice of kernel is not limited to the Gaussian. Precise details of these requirements can be seen in Devroye and Wagner (1980). However, the overall rate of convergence of σ_n for Theorem 2 is dictated by the rate required for the following lemma. The proof of the following lemma also relies on the rotational invariance of the Gaussian kernel. In addition we will also make use of the distribution of the squared norm of the corresponding random variable, which is known in the case of the Gaussian distribution.

Lemma 5. *Let $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ be an i.i.d. sample of realisations of a continuous random variable, $\mathbf{X}$, over $\mathbb{R}^d$, with bounded Lipschitz continuous probability density $p : \mathbb{R}^d \to \mathbb{R}^+$. Let S be a smooth surface which partitions $\mathbb{R}^d$ into S_1 and S_2 such that $0 < \mathbb{P}(\mathbf{X} \in S_1) < 1$. Let $\{\sigma_i\}_{i=1}^{\infty}$ be a null sequence satisfying $n\sigma_n^{2d+2+\epsilon} \to \infty$ for some $\epsilon > 0$. Then,*

$$\frac{c_d \sqrt{2\pi}}{n^2 \sigma_n^{d+1}} \sum_{i:\mathbf{x}_i \in S_1} \sum_{j:\mathbf{x}_j \in S_2} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) \xrightarrow{a.s.} \oint_S p(\mathbf{y})^2 d\mathbf{y} \text{ as } n \to \infty,$$

where $c_d = (\int_{\mathbb{R}^d} k(\|\mathbf{x}\|) d\mathbf{x})^{-1}$.

Lemmas 4 and 5 show immediately that Theorem 2 holds. This is because a ratio of almost surely convergent sequences converges almost surely to the ratio of limits, provided the denominator limit is non-zero. This is ensured in this case by the assumption that $0 < \mathbb{P}(\mathbf{X} \in S_1) < 1$.

Remark 6. *An alternative normalisation of the graph cut, which is referred to as Ratio Cut (Hagen and Kahng, 1992), is given as*

$$\textit{Ratio Cut}(\mathcal{C}, \mathcal{X} \setminus \mathcal{C}) = \textit{Cut}(\mathcal{C}, \mathcal{X} \setminus \mathcal{C}) \left(\frac{1}{|\mathcal{C}|} + \frac{1}{|\mathcal{X} \setminus \mathcal{C}|}\right).$$

Lemma 5 shows almost immediately that the Ratio Cut measured across S, appropriately scaled, converges almost surely to

$$\oint_S p(\mathbf{y})^2 d\mathbf{y} \left(\frac{1}{\mathbb{P}(\mathbf{X} \in S_1)} + \frac{1}{\mathbb{P}(\mathbf{X} \in S_2)}\right),$$

as $n \to \infty$.

3.2 Automating spectral clustering using the asymptotic properties of the normalised cut

Recall that in density clustering clusters are associated with connected regions of high density surrounding the modes of a probability density function. In practice, however, the underlying distribution is unknown and must be estimated. This leads to an appealing interpretation of clusters as regions of high data density, separated by data sparse regions. The practical limitations of density clustering are that it becomes computationally prohibitive to compute these connected high density regions for large and high dimensional datasets, and that density estimation becomes less reliable as dimensionality increases (Rinaldo and Wasserman, 2010; Menardi and Azzalini, 2014).

As discussed in the previous subsection, if the probability distribution underlying the data is multimodal, then the spectral clustering solution will tend to separate clusters by regions of low density, while also assigning high density regions, or more specifically modes, to different clusters. We expect therefore that a clustering solution arising from spectral clustering is likely to satisfy the objectives of density clustering. We can use this observation to determine an appropriate number of clusters to extract using spectral clustering, as follows. Suppose that $\{\mathcal{C}_1, \ldots, \mathcal{C}_c\}$ represents a clustering solution arising from spectral clustering, for a number of clusters c. If c is at most the number of modes of the probability density function then we expect that the cluster boundaries will pass through low density regions and that each cluster will contain at least one mode. On the other hand, if c is greater than the number of modes we either have that a mode is split among multiple clusters or that at least one cluster contains only low density regions. Motivated by this observation, we consider using spectral clustering to propose potential clustering solutions, and assessing their validity in the context of density clustering. Specifically, for a range of values c, we use spectral clustering to obtain a c-way clustering of $\mathcal{X}$. For each cluster we then determine if it is separated from the rest of the clusters only by low density regions. The largest value of c for which all clusters are separated from the remainder then represents the correct number of clusters to extract.

To determine if a cluster, $\mathcal{C} \subset \mathcal{X}$, is separated from the remainder we attempt to establish a path connecting $\mathcal{C}$ to $\mathcal{X} \setminus \mathcal{C}$ which does not pass through any relatively low density regions. If we fail to find such a path then we conclude that the cluster is separated. A point $\mathbf{x} \in \mathbb{R}^d$ has relatively low density if the estimated density at that point satisfies $\hat{p}(\mathbf{x}) < \lambda \min\{\max_{\mathbf{y}\in\mathcal{C}} \hat{p}(\mathbf{y}), \max_{\mathbf{y}\in\mathcal{X}\setminus\mathcal{C}} \hat{p}(\mathbf{y})\}$, where $\lambda \in (0, 1]$ is a chosen setting. We cannot, of course, consider all possible paths, and instead adopt the following heuristic. Noticing that a path connecting a point in $\mathcal{C}$ to a point in $\mathcal{X} \setminus \mathcal{C}$ must pass through the boundary of $\mathcal{C}$, we need only to consider paths connecting the boundary of $\mathcal{C}$ to $\mathcal{X} \setminus \mathcal{C}$. To do this we define the boundary points of $\mathcal{C} \subset \mathcal{X}$ to be those elements which are nearest to some element of $\mathcal{X} \setminus \mathcal{C}$. That is,

$$\text{Boundary}(\mathcal{C}) = \bigcup_{\mathbf{x}\in\mathcal{X}\setminus\mathcal{C}} \arg\min_{\mathbf{y}\in\mathcal{C}} d(\mathbf{x}, \mathbf{y}). \tag{8}$$

Then for each $\mathbf{x} \in \text{Boundary}(\mathcal{C})$ we find the nearest point in $\mathcal{X} \setminus \mathcal{C}$, $\mathbf{y} = \arg\min_{\mathbf{z}\in\mathcal{X}\setminus\mathcal{C}} d(\mathbf{x}, \mathbf{z})$, and then search over the line segment $[\mathbf{x}, \mathbf{y}]$ for points of relatively low density. If there exists such a pair $\mathbf{x}, \mathbf{y}$ with no points of relatively low density between them then we conclude that $\mathcal{C}$ is not separate from the remaining clusters, as there is a path connecting them which avoids low density regions. Pseudocode for the above procedure is given in Algorithm 1.

Remark 7. *Recall that* $\hat{p}(\mathbf{x}_i) \propto \sum_{j=1}^{n} k(\|\mathbf{x}_i - \mathbf{x}_j\|/\sigma) = D_{i,i} + 1$. *We therefore do not have to compute the density values from scratch, except when searching over line segments joing pairs of points in* $\mathcal{X}$. *All pairwise distances, which are used to find the boundary points of* $\mathcal{C}$ *and* $\mathcal{X} \setminus \mathcal{C}$, *are also already computed to obtain the spectral clustering solution.*

A more subtle point arising from Theorem 2, and one which is a caveat concerning the above methodology, may be interpreted in the context of outliers. A dataset containing outliers may be thought of as having arisen from a mixture distribution comprising a main distribution, which contributes to the structure of interest and which represents the vast majority of probability mass; and a long-tailed distribution with low probability mass which produces a small number of outlying points. The length of the tail of this "outlier" distribution dictates how susceptible, in general, an algorithm will be to the outliers. Theorem 2 helps to provide insight into the susceptibility of spectral clustering to outliers. For example, if the tail of the distribution is such that

$$\lim_{b\to\infty} \frac{\oint_{\|\mathbf{y}\|=b} p(\mathbf{y})^2 d\mathbf{y}}{\int_{\|\mathbf{y}\|>b} p(\mathbf{y})^2 d\mathbf{y}} = 0,$$

Algorithm 1: is.density.separated($\mathcal{C}, \mathcal{X} \setminus \mathcal{C}; \lambda$)

$\hat{p}_{\text{thresh}} \leftarrow \min\{\max_{i:\mathbf{x}_i \in \mathcal{C}} D_{i,i} + 1, \max_{i:\mathbf{x}_i \in \mathcal{X}\setminus\mathcal{C}} D_{i,i} + 1\}$
for $\mathbf{x} \in$ Boundary($\mathcal{C}$) **do**
  $\mathbf{y} \leftarrow \arg\min_{\mathbf{z} \in \mathcal{X}\setminus\mathcal{C}} d(\mathbf{x}, \mathbf{z})$
  if $\min_{\gamma \in [0,1]} \sum_{i=1}^{n} k\left(\frac{\|\gamma\mathbf{x} + (1-\gamma)\mathbf{y} - \mathbf{x}_i\|}{\sigma}\right) \geq \lambda \hat{p}_{\text{thresh}}$ **then**
    return FALSE
  end if
end for
return TRUE

then the spectral clustering solution will tend to focus on separating only few points arising in the tail. We would not expect in this case that the density in regions separating these "outlier clusters" from the remainder will be lower than the density *within* these outlier clusters. We identify outlier clusters solely based on their size, and therefore do not check if clusters containing fewer than a prespecified number of points are density separated from the remainder, and only focus on the more substantial clusters. How these outliers are interpreted after the fact is open to a user. If the presence or absence of outliers is in itself a point of interest then these may be inspected separately. We simply merge outliers with their nearest non-outlier cluster, as for the context of this article this makes comparison with other flat clustering algorithms more straightforward.

Finally, recall the requirement on the sequence of scaling parameters given in Theorem 2, that $\lim_{n\to\infty} n\sigma_n^{2d+2+\epsilon} = \infty$ for some $\epsilon > 0$ is sufficient to obtain almost sure convergence of the normalised cut. Although this does not dictate a specific value of the scaling parameter for spectral clustering, it does suggest how σ should decrease as the sample size increases. In particular setting $\sigma = sn^{-1/(2d+3)}$ is appropriate, where s is some stable measure of scale which can be computed from the dataset. The value of s should be almost surely bounded both above and away from zero to ensure convergence of the normalised cut.

Combining the above points we arrive at an algorithm which we call Spectral Partitioning Using Density Separation (SPUDS), pseudocode for which can be found in Algorithm 2. For an initial number, c_0, we compute the spectral clustering solution with c_0 clusters. If all clusters whose sizes are above a chosen size (the non-outlier clusters) are density separated from the remainder, then the best solution must have at least c_0 clusters. We therefore repeatedly increase the number of clusters by one until the spectral clustering solution yields a non-outlier cluster which is not density separated from the remainder. The clustering solution for the largest number of clusters in which all non-outlier clusters are density separated is stored.

On the other hand, if in the clustering solution containing c_0 clusters there is a non-outlier cluster which is not density separated from the remainder, then the correct number of clusters is less than c_0. We therefore repeatedly decrease the number of clusters by one and store the first clustering solution in which all non-outlier clusters are density separated.

Finally, each outlier cluster in the stored clustering solution is merged with its nearest non-outlier cluster, and the resulting clusters are returned.

One practical note on this method is as follows: if either c_0 is much smaller than the final number of clusters, or if outlying points are sufficiently distant that spectral clustering favours splitting off more and more of the tail of the data over producing more substantial clusters, then this method may run slowly. It can be accelerated by, rather than increasing c by one with each iteration when seeking additional clusters, increasing it by a larger number. Then, when a solution is reached in which a non-outlier clusters is not density separated from the remainder, the value of c can be fine tuned, decreasing it repeatedly by one as in the original algorithm.

4 Experimental Results

In this section we will investigate empirically the performance of the algorithm presented in the Section 3. For all experiments we use the following settings. To measure the scale of each dataset we rely on the eigenvalues of its covariance matrix. Although these eigenvalues measure the linear scale of the data, where spectral clustering may be better informed by a more flexible scale measure, they benefit from simplicity

Algorithm 2: Spectral Partitioning Using Density Separation (SPUDS)

Input: Data $\mathcal{X}$, initial cluster number c_0, density path threshold λ, scale measure function s, outlier cluster size threshold γ.

$\sigma \leftarrow s(\mathcal{X})n^{-1/(3+2d)}$

$A \leftarrow A_{i,j} = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}\right), A_{ii} = 0$

$D \leftarrow \text{diag}\left(\sum_{i=1}^n A_{i,1}, \ldots, \sum_{i=1}^n A_{i,n}\right)$

$L \leftarrow I_n - D^{-1/2}AD^{-1/2}$

$c \leftarrow c_0$

for $i \in \{1, \ldots c\}$ **do**

  $\mathbf{u}_i \leftarrow$ eigenvector associated with the i^{th} smallest eigenvalue of L

end for

$U \leftarrow [\mathbf{u}_1, \ldots, \mathbf{u}_c]$

$\{\mathcal{C}_1, \ldots, \mathcal{C}_c\} \leftarrow k\text{-means}(D^{-1/2}U, c)$

if $\forall i \in \{1, \ldots, c\} : |\mathcal{C}_i| < \gamma$ **or** is.density.separated$(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i; \lambda)$ **then**

  repeat

    $\{\mathcal{C}'_1, \ldots, \mathcal{C}'_c\} \leftarrow \{\mathcal{C}_1, \ldots, \mathcal{C}_c\}$

    $c \leftarrow c + 1$

    $\mathbf{u}_c \leftarrow$ eigenvector associated with the c^{th} smallest eigenvalue of L

    $U \leftarrow [\mathbf{u}_1, \ldots, \mathbf{u}_c]$

    $\{\mathcal{C}_1, \ldots, \mathcal{C}_c\} \leftarrow k\text{-means}(D^{-1/2}U, c)$

  until $\exists i \in \{1, \ldots, c\} : |\mathcal{C}_i| \geq \gamma, \neg$is.density.separated$(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i; \lambda)$

  $c \leftarrow c - 1$

  $\{\mathcal{C}_1, \ldots, \mathcal{C}_c\} \leftarrow \{\mathcal{C}'_1, \ldots, \mathcal{C}'_c\}$

else

  repeat

    $c \leftarrow c - 1$

    $U \leftarrow [\mathbf{u}_1, \ldots, \mathbf{u}_c]$

    $\{\mathcal{C}_1, \ldots, \mathcal{C}_c\} \leftarrow k\text{-means}(D^{-1/2}U, c)$

  until $\forall i \in \{1, \ldots, c\} : |\mathcal{C}_i| < \gamma$ **or** is.density.separated$(\mathcal{C}_i, \mathcal{X} \setminus \mathcal{C}_i; \lambda)$

end if

$O \leftarrow \bigcup_{i:|\mathcal{C}_i|<\gamma}\{i\}$

for $i \in O$ **do**

  $i' \leftarrow \arg\min_{j \in \{1,\ldots,c\}\setminus O} d(\mathcal{C}_i, \mathcal{C}_j)$

  $\mathcal{C}_{i'} \leftarrow \mathcal{C}_{i'} \cup \mathcal{C}_i$

end for

return $\bigcup_{i \in \{1,\ldots,c\}\setminus O} \mathcal{C}_i$

and interpretability. In particular we define $s(\mathcal{X}) = \sqrt{\bar{\epsilon}}$, where $\bar{\epsilon}$ is the average of the largest d' eigenvalues of the covariance of $\mathcal{X}$, and d' is an estimate of the intrinsic dimension of $\mathcal{X}$ for which we use Kaiser's criterion (Kaiser, 1960). We found this setting to work fairly consistently for datasets of moderate dimensionality, but for some higher dimensional examples Kaiser's criterion resulted in too small a value of σ. We therefore set d' to 20 if Kaiser's criterion selects a greater value. We distinguish outlier clusters from others as those which contain fewer than $n/200$ points. We set c_0, the initial number of clusters, to 30. Finally we set λ, the density threshold parameter, to 1. This value corresponds to the strictest interpretation that clusters should not contain more than one mode. Note that for lower values of the scaling parameter, σ, the variance of the pointwise density estimates $\hat{p}(\mathbf{x}_i)$ is greater. In such cases it may be beneficial to accommodate some of this variability by setting $\lambda < 1$, thereby potentially connecting pairs of modes in $\hat{p}$ when their heights are not that well distinguished from the density between them.

In addition, we also experimented with the following existing algorithms in order to draw comparisons with the proposed method.

1. Self-Tuning Spectral Clustering (STSC): A fully automatic spectral clustering algorithm which uses a local scaling method to compute similarities and which selects the number of clusters which, through rotation, provides the best alignment of the eigenvectors of the Laplacian matrix with the canonical basis vectors (Zelnik-Manor and Perona, 2004). We used the implementation provided by the authors available from `http://www.vision.caltech.edu/lihi/Demos/SelfTuningClustering.html`.

2. Spectral Clustering using Cluster Distortion (SCCD): Spectral clustering which selects the scaling parameter as that which provides the minimum cluster distortion within the eigenvectors of the Laplacian matrix (Ng et al., 2001). This method cannot automatically determine the number of clusters.

3. Alternative Spectral Clustering (Alt.SC): An automatic spectral clustering method which uses a similar scaling approach as STSC, but which uses strictly smaller scale values, and which selects the number of clusters using an adaptation of Bartlett's test for equal variances. We used the implementation in the `R` package `speccalt` (Bruneau, 2013), available from the `CRAN` repository.

4. k-means: The ubiquitous k-means clustering algorithm is based on minimising the sum of squared distances between data and their assigned cluster's mean. We use the default implementation in `R`, which is based on the algorithm of Hartigan and Wong (1979). To select the number of clusters we used the Gap Statistic (Tibshirani et al., 2001), and set the maximum number of clusters to be twice the actual number.

5. Optimal Extraction of Clusters from Hierarchies (OCE): The algorithm of Campello et al. (2013) applied to cluster hierarchies constructed based on connectivity of clusters using single and average linkage as well as Ward's method. We report only those results based on Ward's method, as this method performed best on the data considered.

6. Gaussian Mixture Model (GMM): We used the method of Fraley and Raftery (2002), which uses the Bayesian Information Criterion to select the number of clusters, and the flexibility of the component covariance matrices.

7. DBSCAN: The method of Ester et al. (1996) is an approximation of the density clustering solution for a kernel density estimate using the uniform kernel. The heuristics for setting the scale and level set parameters described by the authors did not produce sensible results. We therefore generated solutions for a range of settings and report the highest performance for each dataset. This therefore likely overestimates the true performance expected from DBSCAN. In addition we considered more flexible density clustering algorithms, but these were either incapable of handling the datasets due to size, or they did not render sensible solutions for any parameter settings tried.

The following collection of benchmark datasets were used for comparison. Optical recognition of handwritten digits[1] (Opt. Digits), Pen-based recognition of handwritten digits[1] (Pen Digits), Multi-feature digits[1] (M.F. Digits), Statlog landsat satellite[1] (Satellite), Yale faces database B 40×30[2], Phoneme[3], Smartphone based activity recognition[1] (Smartphone), Isolet[1], and Statlog image segmentation[1] (Image Seg.). We evaluate

[1] `https://archive.ics.uci.edu/ml/datasets.html`
[2] `https://cervisia.org/machine_learning_data.php`
[3] `https://statweb.stanford.edu/~tibs/ElemStatLearn/data.html`

clustering solutions using the Normalised Mutual Information (Strehl and Ghosh, 2002), which computes the proportion of shared information in the clustering result and the true clusters.Results from these experiments are summarised in Table 1. We see that the proposed algorithm obtains the highest or close to highest performance in all but one case (Satellite), in which it substantially underestimates the number of clusters. It frequently obtained the highest performance of all algorithms considered, and importantly consistently outperforms other spectral clustering algorithms.

Table 1: Clustering performance based on Normalised Mutual Information (NMI) on benchmark datasets. The highest performance in each case is highlighted in bold.

		SPUDS	STSC	SCCD	Alt.SC	k-means	OCE	GMM	DBSCAN
Opt. Digits	NMI	**0.79**	0.73	0.01	0.01	0.68	0.59	0.63	0.56
(c=10, n=5620, d=64)	c	13	9	-	3	19	4	9	8
Pen Digits	NMI	0.70	0.38	0.19	**0.80**	0.73	0.64	0.73	0.69
(c=10, n=10992, d=16)	c	9	2	-	19	20	5	9	27
M.F. Digits	NMI	0.79	0.71	**0.80**	0.73	0.72	0.57	0.00	0.65
(c=10, n=2000, d=216)	c	18	8	-	8	20	3	1	7
Satellite	NMI	0.40	0.39	0.62	**0.66**	0.60	0.38	0.55	0.51
(c=6, n=6435, d=36)	c	3	2	-	7	11	3	9	5
Yale Faces	NMI	**0.84**	0.04	0.73	0.28	0.70	0.81	0.78	0.54
(c=10, n=5850, d=1200)	c	14	2	-	7	20	10	9	47
Phoneme	NMI	0.69	0.66	**0.87**	0.63	0.68	0.68	0.61	0.37
(c=5, n=4509, d=256)	c	7	3	-	3	10	3	4	3
Smartphone	NMI	0.55*	**0.57**	0.51	0.49	0.56	0.57	0.00	0.41
(c=12, n=10929, d=561)	c	7*	2	-	2	24	2	1	3
Isolet	NMI	**0.72***	0.64	0.70	0.26	0.70	0.42	0.40	0.25
(c=26, n=6238, d=617)	c	25*	15	-	2	52	2	2	5
Image Seg.	NMI	**0.68**	0.42	0.03	0.01	0.61	0.46	0.62	0.50
(c=26, n=2310, d=19)	c	6	3	-	3	14	2	8	5
Average		**0.68**	0.50	0.50	0.43	0.66	0.57	0.48	0.50
Average Regret		**0.06**	0.24	0.25	0.32	0.08	0.18	0.27	0.25
Average Rank		**1.44**	4.06	3.22	4.17	2.28	3.78	3.94	5.11

* indicates that the accelerated version of the algorithm, described at the end of Section 3, was used. The number of clusters was increased by 10 with each iteration.

In addition to the performance on the individual datasets, we also report the average performance over all datasets, as well as the average regret and average rank. The regret of an algorithm on a dataset is the difference between the best performing method on that dataset and the performance of the algorithm. The rank of an algorithm on a dataset is its position in the ordered set of performance values on that dataset. It is clear that overall the proposed algorithm is competitive with existing algorithms. Again the comparisons with existing spectral clustering algorithms is compelling. The only algorithm with similar performance overall is k-means. Note however that because of the high computation time associated with estimating the number of clusters via the Gap Statistic, we enforced an upper bound on the number of clusters which could be extracted. In most cases this upper bound was selected by the Gap. It is likely therefore that without such a reasonable upper bound this approach would even more significantly overestimate the number of clusters, thereby causing performance to deteriotate.

5 Discussion

In this paper the asymptotic value of the normalised cut was investigated, and a new spectral clustering algorithm proposed which makes use of this asymptotic result. Speficically, based on the asymptotic value of the normalised cut we expect that clustering solutions arising from spectral clustering will both separate clusters by regions of low empirical density, and also assign high density regions to each cluster. This observation is used to automatically determine the number of clusters, as the largest value for which all clusters remain separated from the remainder by low density regions. Furthermore requirements on the scaling parameter to ensure convergence of the normalised cut are used to provide a data driven method for

setting this important parameter. Experiments on a collection of benchmark datasets indicate that this new algorithm is competitive with state of the art clustering methods.

Appendix: Proofs of Lemmas 4 and 5

Throughout the remainder we will assume that the kernel k is normalised so that $\int_{\mathbb{R}^d} k(\|\mathbf{x}\|) = 1$. Notice that this does not affect the value of the normalised cut, as the normalisation is applied in both the numerator and denominator terms.

Lemma 4

Let $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ be an i.i.d. sample of realisations of a continuous random variable, $\mathbf{X}$, over $\mathbb{R}^d$, with Lipschitz continuous probability density $p : \mathbb{R}^d \to \mathbb{R}^+$. Let $\{\sigma_i\}_{i=1}^{\infty}$ be a null sequence satisfying, $n\sigma_n^{2d+2+\epsilon} \to \infty$ for some $\epsilon > 0$. Let $M \subset \mathbb{R}^d$ be measurable. Then,

$$\frac{c_d}{n^2\sigma_n^d} \sum_{i:\mathbf{x}_i \in M} \sum_{j \neq i}^{n} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) \xrightarrow{a.s.} \int_M p(\mathbf{x})^2 d\mathbf{x}, \text{ as } n \to \infty,$$

where $c_d = (\int_{\mathbb{R}^d} k(\|\mathbf{x}\|)d\mathbf{x})^{-1}$.

Proof. First notice that

$$\frac{1}{n^2\sigma_n^d} \sum_{i:\mathbf{x}_i \in M} \sum_{j \neq i} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) = \frac{1}{n^2\sigma_n^d} \sum_{i:\mathbf{x}_i \in M} \sum_{j=1}^{n} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) + \mathcal{O}(n^{-1}\sigma_n^{-d})$$

Now,

$$\frac{1}{n^2\sigma_n^d} \sum_{i:\mathbf{x}_i \in M} \sum_{j=1}^{n} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) = \frac{1}{n} \sum_{i:\mathbf{x}_i \in M}^{n} p(\mathbf{x}_i) + \frac{1}{n} \sum_{i:\mathbf{x}_i \in M} \left(\frac{1}{n\sigma_n^d} \sum_{j=1}^{n} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) - p(\mathbf{x}_i)\right).$$

The first term is a standard Monte Carlo estimate of the integral $\int_M p(\mathbf{x})^2 d\mathbf{x}$ and satisfies,

$$\left|\frac{1}{n} \sum_{i:\mathbf{x}_i \in M}^{n} p(\mathbf{x}_i) - \int_M p(\mathbf{x})^2 d\mathbf{x}\right| \xrightarrow{a.s.} 0 \text{ as } n \to \infty.$$

The second term represents the average error of a kernel estimator of p, and clearly satisfies,

$$\left|\frac{1}{n} \sum_{i:\mathbf{x}_i \in M} \left(\frac{1}{n\sigma_n^d} \sum_{j=1}^{n} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) - p(\mathbf{x}_i)\right)\right| \leq \sup_{\mathbf{x} \in \mathbb{R}^d} \left|\frac{1}{n\sigma_n^d} \sum_{j=1}^{n} k\left(\frac{\|\mathbf{x} - \mathbf{x}_j\|}{\sigma_n}\right) - p(\mathbf{x})\right|,$$

where $\frac{1}{n\sigma_n^d} \sum_{j=1}^{n} k(\|\mathbf{x} - \mathbf{x}_j\|/\sigma_n)$ is the kernel based estimate of $p(\mathbf{x})$. Now, since p is Lipschitz continuous, and since the Gaussian kernel clearly satisfies the conditions of (Devroye and Wagner, 1980, Theorem 1) for $\sigma_n \to 0, n\sigma_n^{2d+2+\epsilon} \to \infty$ for any $\epsilon > 0$, the kernel estimate converges uniformly almost surely. Specifically,

$$\sup_{\mathbf{x} \in \mathbb{R}^d} \left|\frac{1}{n\sigma_n^d} \sum_{j=1}^{n} k\left(\frac{\|\mathbf{x} - \mathbf{x}_j\|}{\sigma_n}\right) - p(\mathbf{x})\right| \xrightarrow{a.s.} 0 \text{ as } n \to \infty.$$

Putting these together gives the result. □

Lemma 5

Let $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ be an i.i.d. sample of realisations of a continuous random variable, $\mathbf{X}$, over $\mathbb{R}^d$, with bounded Lipschitz continuous probability density $p : \mathbb{R}^d \to \mathbb{R}^+$. Let S be a smooth surface which partitions $\mathbb{R}^d$ into S_1 and S_2 such that $0 < \mathbb{P}(\mathbf{X} \in S_1) < 1$. Let $\{\sigma_i\}_{i=1}^{\infty}$ be a null sequence satisfying $n\sigma_n^{2d+2+\epsilon} \to \infty$ for some $\epsilon > 0$. Then,

$$\frac{c_d\sqrt{2\pi}}{n^2\sigma_n^{d+1}} \sum_{i:\mathbf{x}_i \in S_1} \sum_{j:\mathbf{x}_j \in S_2} k\left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma_n}\right) \xrightarrow{a.s.} \oint_S p(\mathbf{y})^2 d\mathbf{y} \text{ as } n \to \infty,$$

where $c_d = (\int_{\mathbb{R}^d} k(\|\mathbf{x}\|)d\mathbf{x})^{-1}$.

Proof. To begin with, take $i \in \{1, \ldots, n\}$, and let $\mathcal{X}' = (\mathcal{X} \setminus \{\mathbf{x}_i\}) \cup \{\mathbf{x}_i'\}$ for any $\mathbf{x}_i' \in \mathbb{R}^d$.

$$\left| \sum_{\mathbf{x}\in\mathcal{X}\cap S_1} \sum_{\mathbf{y}\in\mathcal{X}\cap S_2} k\left(\frac{\|\mathbf{x}-\mathbf{y}\|}{\sigma_n}\right) - \sum_{\mathbf{x}\in\mathcal{X}'\cap S_1} \sum_{\mathbf{y}\in\mathcal{X}'\cap S_2} k\left(\frac{\|\mathbf{x}-\mathbf{y}\|}{\sigma_n}\right) \right| \leq n,$$

and hence the random variable $\nu = \frac{\sqrt{2\pi}}{n^2\sigma_n^{d+1}} \sum_{i:\mathbf{x}_i\in S_1} \sum_{j:\mathbf{x}_j\in S_2} k\left(\frac{\|\mathbf{x}_i-\mathbf{x}_j\|}{\sigma_n}\right)$ is difference bounded by $\frac{\sqrt{2\pi}}{n\sigma_n^{d+1}}$. McDiarmid's inequality therefore ensures,

$$\mathbb{P}\left(|\nu - \mathbb{E}[\nu]| > \epsilon\right) < \exp\left(-\frac{2\epsilon^2 n\sigma_n^{2d+2}}{2\pi}\right).$$

By assumption on the rate of convergence of σ_n to zero, it is therefore sufficient to show that $\mathbb{E}[\nu] \to \oint_S p(\mathbf{y})^2 d\mathbf{y}$ as $\sigma_n \to 0$. We henceforth drop the notational dependence on n and write only σ.

Establishing the convergence of $\mathbb{E}[\nu]$ will be achieved in a number of steps. To begin with consider that,

$$\mathbb{E}[\nu] = \frac{\sqrt{2\pi}}{\sigma}\mathbb{P}(\mathbf{X} \in S_1)\mathbb{P}(\mathbf{X} \in S_2)\mathbb{E}\left[\frac{1}{\sigma^d}k\left(\frac{\|\mathbf{Y}-\mathbf{Z}\|}{\sigma}\right)\right],$$

where $\mathbf{Y}$ and $\mathbf{Z}$ are random variables representing the truncations of $\mathbf{X}$ in S_1 and S_2 respectively. Therefore,

$$\mathbb{E}\left[\frac{1}{\sigma^d}k\left(\frac{\|\mathbf{Y}-\mathbf{Z}\|}{\sigma}\right)\right] = \int_{S_1}\int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) \frac{p(\mathbf{y})}{\mathbb{P}(\mathbf{X}\in S_1)} \frac{p(\mathbf{z})}{\mathbb{P}(\mathbf{X}\in S_2)} d\mathbf{z}d\mathbf{y}$$

$$\Rightarrow \mathbb{E}[\nu] = \frac{\sqrt{2\pi}}{\sigma}\int_{S_1}\int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) p(\mathbf{y})p(\mathbf{z})d\mathbf{z}d\mathbf{y}.$$

Next we use the Lipschitz continuity of p to show that the product $p(\mathbf{y})p(\mathbf{z})$ can be replaced by $p(\mathbf{y})^2$ while inducing very small error. This is because the kernel weight $k(\|\mathbf{y}-\mathbf{z}\|/\sigma)$ is only large when $\mathbf{z}$ is close to $\mathbf{y}$, and hence $p(\mathbf{z}) \approx p(\mathbf{y})$. Suppose p has Lipschitz constant L. Then we have

$$p(\mathbf{y})\int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right)(p(\mathbf{y}) - L\|\mathbf{z}-\mathbf{y}\|)d\mathbf{z} \leq \int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) p(\mathbf{y})p(\mathbf{z})d\mathbf{z}$$

$$\leq p(\mathbf{y})\int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right)(p(\mathbf{y}) + L\|\mathbf{z}-\mathbf{y}\|)d\mathbf{z}.$$

Now, if $\mathbf{y}$ is s.t. $d(\mathbf{y}, S) \geq \sqrt{\sigma}$, then for small σ we have

$$\int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right)\|\mathbf{y}-\mathbf{z}\|d\mathbf{z} \leq \int_{\mathbf{z}:\|\mathbf{z}-\mathbf{y}\|\geq\sqrt{\sigma}} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right)\|\mathbf{y}-\mathbf{z}\|d\mathbf{z}$$

$$= \mathcal{O}\left(\sqrt{\sigma}\left(1 - \Phi(\sigma^{-0.5})\right)\right) = \mathcal{O}(\sigma^2),$$

where Φ is the distribution function of the standard Gaussian random variable. It is clear that in fact this quantity is of a far smaller order of magnitude than σ^2 as σ approaches zero, however this is sufficient for our purposes. On the other hand, if $d(\mathbf{y}, S) < \sqrt{\sigma}$, then we instead have

$$\int_{S_2} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right)\|\mathbf{y}-\mathbf{z}\|d\mathbf{z} \leq \int_{\mathbb{R}^d} \frac{1}{\sigma^d}k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right)\|\mathbf{y}-\mathbf{z}\|d\mathbf{z}$$

$$= \mathcal{O}(\sigma),$$

as the latter integral is the expected distance of a standard d dimensional Gaussian random variable from its mean, and is $\mathcal{O}(\sqrt{d}\sigma) = \mathcal{O}(\sigma)$.

Next, using a similar geometric construction to that in (Narayanan et al., 2006), let r be the radius of the largest ball which can be placed tangent to S at any point and which intersects S at a single point. Then, for $\mathbf{y} \in S_1$ s.t. $d(\mathbf{y}, S) < \min\{r, \sqrt{\sigma}\}$ let $\mathbf{y}'$ be the nearest point to $\mathbf{y}$ lying on S ($\mathbf{y}'$ is unique since $d(\mathbf{y}, S) < r$). Define H to be the hyperplane tangent to S at $\mathbf{y}'$ and let H_1 be the corresponding half space containing $\mathbf{y}$ and H_2 the half space not containing $\mathbf{y}$. We will show that as σ approaches zero,

$$\int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} = \int_{H_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} + \mathcal{O}(\sigma).$$

To that end, let B_1 be the ball of radius r tangent to S at $\mathbf{y}'$ and lying in S_1. Define B_2 similarly except $B_2 \subset S_2$. We clearly have

$$\left| \int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} - \int_{H_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} \right| \leq \int_{\mathbb{R}^d \setminus B_1} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} - \int_{B_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z}.$$

By translation so that the centre of B_1 lies at the origin, we have

$$\int_{\mathbb{R}^d \setminus B_1} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} = 1 - \mathbb{P}\left(\|\mathbf{W}\|^2 \leq r^2\right),$$

where $\mathbf{W} \sim \mathcal{N}(\mathbf{y}, \sigma^2 I)$. Therefore, the random variable $\|\mathbf{W}\|^2/\sigma^2$ is distributed χ^2 with d degrees of freedom and non-centrality parameter $\frac{1}{\sigma^2}(r - d(\mathbf{y}, S))^2$. For brevity let $\eta = d(\mathbf{y}, S)$. We therefore have,

$$\int_{\mathbb{R}^d \setminus B_1} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} = Q_{d/2}\left(\frac{r-\eta}{\sigma}, \frac{r}{\sigma}\right),$$

where $Q_\nu(\alpha, \beta)$ is the Marcum Q-function. Similarly, we find that

$$\int_{B_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} = 1 - Q_{d/2}\left(\frac{r+\eta}{\sigma}, \frac{r}{\sigma}\right).$$

Now, for $z \in \mathbb{Z}$, it has been shown that (Sun and Zhou, 2008)

$$Q_{z+0.5}(\alpha, \beta) = \frac{1}{2}\text{erfc}\left(\frac{\alpha+\beta}{\sqrt{2}}\right) + \frac{1}{2}\text{erfc}\left(\frac{\beta-\alpha}{\sqrt{2}}\right) + \left(\frac{1}{2\pi\alpha\beta}\right)^{0.5} \exp\left(-\frac{\alpha^2+\beta^2}{2}\right) \sum_{k=0}^{z-1} \left(\frac{\beta}{\alpha}\right)^{k+0.5} \times \sum_{r=0}^{k} \frac{(k+r)!}{r!(k-r)!(2\alpha\beta)^r} \left((-1)^r \exp(\alpha\beta) + (-1)^{k+1} \exp(-\alpha\beta)\right),$$

where erfc is the complementary error function $\text{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty \exp(-z^2) dz$. With some straightforward algebra, and for $\alpha, \beta \gg 0$ we can arrive at,

$$Q_{z+0.5}(\alpha, \beta) = \frac{1}{2}\text{erfc}\left(\frac{\alpha+\beta}{\sqrt{2}}\right) + \frac{1}{2}\text{erfc}\left(\frac{\beta-\alpha}{\sqrt{2}}\right) + \mathcal{O}(\beta^{z-1}\alpha^{-z}), \qquad \beta \geq \alpha$$

$$Q_{z+0.5}(\alpha, \beta) = \frac{1}{2}\text{erfc}\left(\frac{\alpha+\beta}{\sqrt{2}}\right) + \frac{1}{2}\text{erfc}\left(\frac{\beta-\alpha}{\sqrt{2}}\right) + \mathcal{O}(\alpha\beta)^{-0.5}, \qquad \text{otherwise.}$$

For d odd we therefore have,

$$\int_{\mathbb{R}^d \setminus B_1} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} - \int_{B_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} \leq \frac{1}{2}\left(\text{erfc}\left(\frac{2r-\eta}{\sqrt{2}\sigma}\right) + \text{erfc}\left(\frac{\eta}{\sqrt{2}\sigma}\right)\right) + \frac{1}{2}\left(\text{erfc}\left(\frac{2r+\eta}{\sqrt{2}\sigma}\right) + \text{erfc}\left(\frac{-\eta}{\sqrt{2}\sigma}\right)\right) - 1 + \mathcal{O}(\sigma).$$

Now, $\operatorname{erfc}(-x)+\operatorname{erfc}(x) = 2$. Therefore,

$$\int_{\mathbb{R}^d \setminus B_1} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} - \int_{B_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{z}-\mathbf{y}\|}{\sigma}\right) d\mathbf{z} \leq \frac{1}{2}\left(\operatorname{erfc}\left(\frac{2r-\eta}{\sqrt{2}\sigma}\right) + \operatorname{erfc}\left(\frac{2r+\eta}{\sqrt{2}\sigma}\right)\right) + \mathcal{O}(\sigma)$$
$$= \mathcal{O}(\sigma)$$

as $\sigma \to 0$, since $r > \eta$ and using the fact that $\operatorname{erfc}(x) \leq \exp(-x)$ for $x > 0$. For d even we can simply use the above and the fact that $Q_\nu(\alpha, \beta)$ is increasing in ν for all fixed α, β. Specifically, $Q_{z-0.5}(\alpha, \beta) < Q_z(\alpha, \beta) < Q_{z+0.5}(\alpha, \beta)$ for $z \in \mathbb{Z}$.

We therefore have shown that, for $\mathbf{y}$ s.t. $d(\mathbf{y}, S) < \min\{r, \sqrt{\sigma}\}$,

$$\left| \int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} - \int_{H_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} \right| = \mathcal{O}(\sigma).$$

On the other hand, if $d(\mathbf{y}, S) \geq \min\{r, \sqrt{\sigma}\}$, then we instead have

$$\left| \int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} - \int_{H_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} \right| = \mathcal{O}(\sigma^2),$$

as $\sigma \to 0$ since both terms in the left hand side are $\mathcal{O}(\sigma^m)$ for all $m \in \mathbb{R}$. The exponent of σ is not crucial for the final result except that it must be strictly greater than 1 in this instance.

The next step simply relies on the recognition that, since k is the Gaussian kernel and the Gaussian distribution is closed under marginalisation, we have

$$\int_{H_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} = 1 - \Phi\left(\frac{\eta}{\sigma}\right).$$

This is achieved by a transformation such that $\mathbf{y}$ lies at the origin and H is defined by the equation $\mathbf{z}_1 = \eta$, and through marginalising out all other dimensions.

Bringing these approximations together, we therefore have for small σ,

$$\begin{aligned}
\mathbb{E}[\nu] &= \frac{\sqrt{2\pi}}{\sigma} \int_{S_1} \int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) p(\mathbf{y}) p(\mathbf{z}) d\mathbf{z} d\mathbf{y} \\
&= \frac{\sqrt{2\pi}}{\sigma} \int_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)<\sqrt{\sigma}}} p(\mathbf{y}) \left(p(\mathbf{y}) \int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} + \mathcal{O}(\sigma) \right) d\mathbf{y} \\
&\quad + \frac{\sqrt{2\pi}}{\sigma} \int_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)\geq\sqrt{\sigma}}} p(\mathbf{y}) \left(p(\mathbf{y}) \int_{S_2} \frac{1}{\sigma^d} k\left(\frac{\|\mathbf{y}-\mathbf{z}\|}{\sigma}\right) d\mathbf{z} + \mathcal{O}(\sigma^2) \right) d\mathbf{y} \\
&= \frac{\sqrt{2\pi}}{\sigma} \int_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)<\sqrt{\sigma}}} p(\mathbf{y}) \left(p(\mathbf{y}) \left(1 - \Phi\left(\frac{d(\mathbf{y},S)}{\sigma}\right) + \mathcal{O}(\sigma) \right) + \mathcal{O}(\sigma) \right) d\mathbf{y} \\
&\quad + \frac{\sqrt{2\pi}}{\sigma} \int_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)\geq\sqrt{\sigma}}} p(\mathbf{y}) \left(p(\mathbf{y}) \left(1 - \Phi\left(\frac{d(\mathbf{y},S)}{\sigma}\right) + \mathcal{O}(\sigma^2) \right) + \mathcal{O}(\sigma^2) \right) d\mathbf{y} \\
&= \frac{\sqrt{2\pi}}{\sigma} \int_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)<\sqrt{\sigma}}} \left(p(\mathbf{y})^2 \left(1 - \Phi\left(\frac{d(\mathbf{y},S)}{\sigma}\right)\right) + p(\mathbf{y})(1+p(\mathbf{y}))\mathcal{O}(\sigma) \right) d\mathbf{y} \\
&\quad + \frac{\sqrt{2\pi}}{\sigma} \int_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)\geq\sqrt{\sigma}}} \left(p(\mathbf{y})^2 \left(1 - \Phi\left(\frac{d(\mathbf{y},S)}{\sigma}\right)\right) + p(\mathbf{y})(1+p(\mathbf{y}))\mathcal{O}(\sigma^2) \right) d\mathbf{y} \\
&= \frac{\sqrt{2\pi}}{\sigma} \int_0^{\sqrt{\sigma}} \oint_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)=\eta}} \left(p(\mathbf{y})^2 \left(1 - \Phi\left(\frac{\eta}{\sigma}\right)\right) + p(\mathbf{y})(1+p(\mathbf{y}))\mathcal{O}(\sigma) \right) d\mathbf{y} d\eta \\
&\quad + \frac{\sqrt{2\pi}}{\sigma} \int_{\sqrt{\sigma}}^{\infty} \oint_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)=\eta}} \left(p(\mathbf{y})^2 \left(1 - \Phi\left(\frac{\eta}{\sigma}\right)\right) + p(\mathbf{y})(1+p(\mathbf{y}))\mathcal{O}(\sigma^2) \right) d\mathbf{y} d\eta \\
&= \frac{\sqrt{2\pi}}{\sigma} \int_0^{\infty} \oint_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)=\eta}} p(\mathbf{y})^2 \left(1 - \Phi\left(\frac{\eta}{\sigma}\right)\right) d\mathbf{y} d\eta + \mathcal{O}(\sqrt{\sigma}).
\end{aligned}$$

Now, since p is bounded and Lipschitz, we know that

$$f : \mathbb{R}^+ \to \mathbb{R}^+ : \eta \mapsto \oint_{\substack{\mathbf{y}:\mathbf{y}\in S_1 \\ d(\mathbf{y},S)=\eta}} p(\mathbf{y})^2 d\mathbf{y}$$

is Lipschitz. Therefore,

$$\begin{aligned}
\mathbb{E}[\nu] &= \frac{\sqrt{2\pi}}{\sigma} \int_0^\infty \left(1 - \Phi\left(\frac{\eta}{\sigma}\right)\right) \left(\oint_S p(\mathbf{y})^2 d\mathbf{y} + \mathcal{O}(\eta)\right) d\eta + \mathcal{O}(\sqrt{\sigma}) \\
&= \oint_S p(\mathbf{y})^2 d\mathbf{y} \left(\frac{\sqrt{2\pi}}{\sigma} \int_0^\infty \left(1 - \Phi\left(\frac{\eta}{\sigma}\right)\right)\right) + \frac{\sqrt{2\pi}}{\sigma} \int_0^\infty \left(1 - \Phi\left(\frac{\eta}{\sigma}\right)\right) \mathcal{O}(\eta) d\eta + \mathcal{O}(\sqrt{\sigma}) \\
&= \oint_S p(\mathbf{y})^2 d\mathbf{y} + \mathcal{O}(\sqrt{\sigma}),
\end{aligned}$$

since $\frac{\sqrt{2\pi}}{\sigma} \int_0^\infty (1 - \Phi(\eta/\sigma)) d\eta = 1$ and $\frac{1}{\sigma} \int_0^\infty (1 - \Phi(\eta/\sigma)) \eta d\eta = \mathcal{O}(\sigma)$. This proves the result. □

References

Bruneau, P. (2013). *speccalt: Alternative spectral clustering, with automatic estimation of k.* R package version 0.1.1.

Campello, R. J., D. Moulavi, A. Zimek, and J. Sander (2013). A framework for semi-supervised and unsupervised optimal extraction of clusters from hierarchies. *Data Mining and Knowledge Discovery 27*(3), 344–371.

Devroye, L. P. and T. J. Wagner (1980). The strong uniform consistency of kernel density estimates. In *Multivariate Analysis V: Proceedings of the fifth International Symposium on Multivariate Analysis*, Volume 5, pp. 59–77.

Ester, M., H.-P. Kriegel, J. Sander, X. Xu, et al. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, Volume 96, pp. 226–231.

Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American statistical Association 97*(458), 611–631.

Hagen, L. and A. B. Kahng (1992). New spectral methods for ratio cut partitioning and clustering. *Computer-aided design of integrated circuits and systems, ieee transactions on 11*(9), 1074–1085.

Hartigan, J. A. and M. A. Wong (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics) 28*(1), 100–108.

Kaiser, H. F. (1960). The application of electronic computers to factor analysis. *Educational and psychological measurement 20*(1), 141–151.

Menardi, G. and A. Azzalini (2014). An advancement in clustering via nonparametric density estimation. *Statistics and Computing 24*(5), 753–767.

Narayanan, H., M. Belkin, and P. Niyogi (2006). On the relation between low density separation, spectral clustering and graph cuts. In *Advances in Neural Information Processing Systems*, pp. 1025–1032.

Ng, A. Y., M. I. Jordan, Y. Weiss, et al. (2001). On spectral clustering: Analysis and an algorithm. In *NIPS*, Volume 14, pp. 849–856.

Rinaldo, A. and L. Wasserman (2010). Generalized density clustering. *The Annals of Statistics*, 2678–2722.

Shi, J. and J. Malik (2000). Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 22*(8), 888–905.

Strehl, A. and J. Ghosh (2002). Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *Journal of machine learning research 3*(Dec), 583–617.

Sun, Y. and S. Zhou (2008). Tight bounds of the generalized marcum q-function based on log-concavity. In *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, pp. 1–5. IEEE.

Tibshirani, R., G. Walther, and T. Hastie (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63*(2), 411–423.

Trillos, N. G., D. Slepcev, J. von Brecht, T. Laurent, and X. Bresson (2015). Consistency of cheeger and ratio graph cuts. *Journal of Machine Learning Research*.

von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and Computing 17*(4), 395–416.

Wagner, D. and F. Wagner (1993). *Between min cut and graph bisection*. Springer.

Zelnik-Manor, L. and P. Perona (2004). Self-tuning spectral clustering. In *Advances in neural information processing systems*, pp. 1601–1608.